EMULEX®

White Paper

Virtual Server-SAN Connectivity: Why blade server virtualization and Fibre Channel SAN connectivity are rapidly being adopted

AT A GLANCE

Server virtualization is rapidly gaining market acceptance for server consolidation and provisioning. The introduction of Emulex LightPulse® Virtual HBA technology, implementing standards-based N-Port ID virtualization, will allow users to combine and leverage virtual servers and storage networks into flexible, secure and reliable solutions.

PRODUCTS

Emulex LightPulse HBAs

LightPulse Virtual HBA Technology

APPLICATIONS

Server virtualization

Enterprise-class SAN connectivity solutions for blades

Server Virtualization and SANs

Server virtualization provides the ability to deploy a discrete number of full-featured application environments called "virtual machines" or "guests" on a single, or a few, hardware platforms. Hardware resource scheduling and management is run behind the scenes by a "hypervisor"– a user-invisible operating system. Virtualization delivers benefits that meet important user needs:

- Server consolidation, yielding an increase in server hardware utilization. Analyst research confirms that server utilization, typically as low as 20% for separate servers, can increase to as much as 75 to 80% on a virtual server. This reduces the overall number of servers required for an equivalent workload and enhances return on server investment.
- Reduced management and maintenance overhead, in direct relation to the reduced number of hardware platforms (freeing up the IT staff for more productive tasks such as new application deployment).
- Cost effective test bed deployment of new applications.
- Faster provisioning for new applications, which only require setting up a new "virtual machine" using available server, network and storage resources.
- Easier incident resumption and disaster recovery, as the "virtual machines" can be restarted faster and more reliably on alternate hardware resources.

These benefits explain why analyst firms such as Gartner Group have rated server virtualization as one of three "disruptive technologies" that will change the face of IT, at par with blade servers and Linux.

Two different breeds of virtual server environments have emerged. Virtual servers were initially introduced on very large mainframes able to

Chart 1: Server Farms

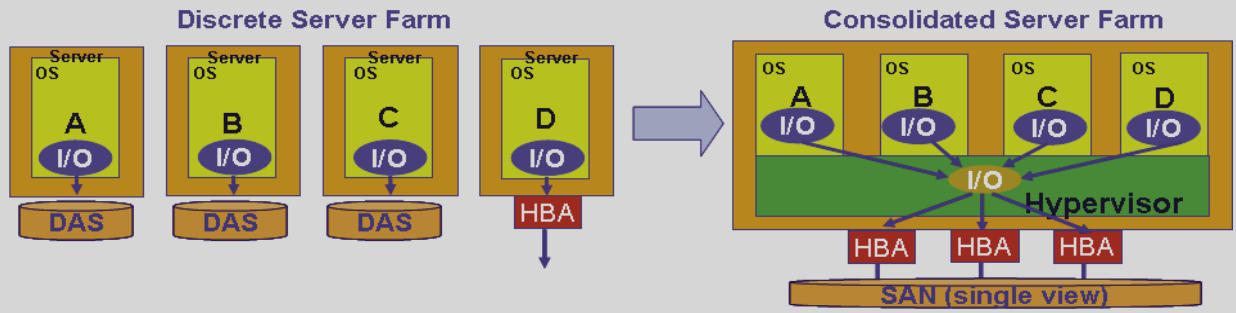

marshal extremely powerful multiprocessor, memory and I/O resources. In such an environment, virtualization creates a way to provide "smaller chunks" of servers for specialized or transient usage. IBM mainframes (eServer zSeries) pioneered server virtualization several decades ago.

A different, increasingly compelling model consolidates Windows and Linux environments onto farms of standard Intel-based servers. The benefits here include resource management, flexibility and workload portability. The leader in this market is VMware, an EMC subsidiary whose ESX environment is being aggressively deployed by Fortune 500 companies.

Server Virtualization on Blades:

The early generations of blade servers offered little synergy with virtualization, as they provided only limited computational and connectivity capabilities under each system image (or server blade), and deployed Linux or Windows management environments that did not account for virtual servers. Several factors are accelerating convergence of blades and virtualization:

- Greater computational power on each blade, especially using the latest multicore technologies
- Innovation in consolidating I/O resources first with N-Port ID Virtualization, and in the future with the evolving IOV standards developed by the PCI SIG
- Coexistence then convergence of the blade and virtualization management environments within the context of DMTF and SNIA standards.
- Virtual appliances (dedicated, pre-configured virtual servers) can be easily and safely implemented on server blades. As ISVs develop their offerings of such appliances, they will be contributing to extend the market for blade-based server virtualization.

Vendors and users will increasingly combine the benefits of physical consolidation (via blade servers) and logical consolidation (via virtualization).

Server Virtualization and the SAN

Consolidation servers, and virtual servers, are mostly located in the main data center and draw on enterprise storage resources. Virtual operating environments such as VMware already enjoy a very high SAN attach rate. The portability and recovery capabilities of these environments rely on external shared storage and are most effective in a SAN connectivity environment. In addition, the high performance delivered by the Fibre Channel protocol is best positioned to serve the higher I/O requirements for the multiple virtual machines running on a single server. SAN connectivity helps enable server virtualization, while server virtualization drives an increased need for SAN connectivity. Today's server virtualization solutions, combined with emerging standards in the storage management space, will bring further success to this growing market segment. N-Port ID Virtualization, in particular, will provide the following incremental benefits:

- The ability to utilize today's proven SAN management methods to associate a SAN-based storage device with a particular virtual machine. This allows the end user to manage this critical aspect of storage management with a common methodology across bare-metal and virtual servers. In cases requiring virtual machine migration,

this feature will also guarantee that access privileges travel with the virtual machine, preventing today's requirement to expose the same storage device to multiple virtual servers.

- The ability to provide within-fabric isolation among not only virtual servers, but also among virtual machines. This isolation will allow network events and device-specific behavior to be confined to the particular virtual machine. This takes the same best practices in place in data center SANs today and extends their reach into the virtual server.
- Application-transparent HBA replacement, making provisioning easier and reducing long-term cost of ownership.

Still, adoption of virtual server environments is somewhat inhibited by the limited flexibility with respect to storage access. All virtual machines on a physical server currently share the same external storage resources, and they share a single set of Fibre Channel addresses provided via the underlying OS or hypervisor. This has several limitations:

- The absence of fabric zoning or array-level LUN masking keeps all storage potentially visible to all virtual machines, contradicting SAN best practices. The current remedy is to implement masking within the server itself, by way of the hypervisor's file system. The limitations of this solution are obvious. Firstly, it introduces another, server administrator-managed, procedure and tool beside those in use by the storage administrator. Secondly, it scales very poorly: rezoning for hundreds of servers only requires one central change, whereas server-level masking must be modified for each virtual server (in addition to zoning for traditional servers).
- Specific SAN resources cannot be assigned by priority to mission-critical applications or virtual machines, which limits quality-of-service (QoS) implementations.
- One of the key features of virtual environments, the ability to quickly overcome hardware failures by moving virtual machines to alternate servers, is restricted and in some cases inhibited by the hardware-based Fibre Channel port addresses being lost along with the failed servers. Here again, users eager to enjoy the benefits of virtualization have a workaround, which is to implement open zoning for all virtual servers. This solution, applicable when the role of the virtual server is limited to non-critical file and print server consolidation, cannot be deployed for mission-critical workloads where zones are used to enforce application and server isolation.

The fact that server virtualization is highly successful under its current storage-attach restrictions testifies to the benefits and potential of this technology. Removing these limits will enable new classes of applications, fuller data center integration, and accelerated virtual server deployment. This will be accomplished with another technology now coming of age: N-Port ID Virtualization (NPIV) .

N-Port ID Virtualization Overview

NPIV is an ANSI T11 standard describing how a single Fibre Channel HBA port can register with the fabric using several worldwide port names (WWPNs). Upon bringup, the HBA first logs into the fabric (FLOGI), then requests (through ADISC) as many fabric IDs from the switch as it has logical ports, identified by a unique WWPN. The fabric will then route messages, display topologies and monitor status for these virtual end points as it already does for the physical ones. The use of multiple addresses through a single HBA port is very valuable in virtual servers as it enables zoning and LUN masking, giving each virtual machine specialized access to only its required storage resources. In addition, exclusive assignment of buffer resources to priority virtual machines, through their "port", provides added granularity to fulfill SLAs. Finally, the ability to "tear down" a virtual port and reinitiate it on a different server greatly enhances virtual machine portability for load balancing, portability and incident recovery. In short, NPIV enhances SAN connectivity, flexibility, resource allocation and recovery.

NPIV requires two capabilities:

- The HBA's firmware and the proximate switch must be able to negotiate several addresses. This is specified as part of the ANSI T11 Fibre Channel standards, and is available on the latest fabric switches from leading vendors such as Brocade, Cisco and McDATA.
- The driver must be able to accept requests from the operating system (in the case of a virtual server, the hypervisor) to create, modify, delete Fibre Channel addresses, and to provide status on each address (status, health, throughput, event history). Emulex is partnering with leading virtual OS providers in defining and implementing the operating system and administrative environment for each technical environment and usage scenario, using a common programming interface.

Emulex Leadership in Virtual Server Connectivity

Emulex has taken a leadership position in connecting virtual servers to the fabric. Emulex is exclusively providing, through our partnership with IBM, adapters supporting NPIV on IBM's zSeries mainframe computers, that are now part of an end-to-end offering deployed in production environments.

Emulex has also taken the lead in providing the same benefit to the volume server virtualization marketplace:

- Emulex and VMware are partnering on virtual-machine-level port creation, management and mobility. Joint demonstrations have been presented in several venues including Storage Networking World and VMworld. NPIV provides the core of a multi-vendor "VMware Community Source" project providing the core capability for upcoming productization.
- Emulex recently announced VMPilot™, a new management software application that works with Emulex LightPulse HBAs to create a virtual HBA connection for Microsoft Virtual Server users. Key VMPilot features include step-by-step wizards to create and migrate virtual machines with virtual HBA SAN connections.
- Emulex has developed and demonstrated the NPIV capability in the Xen open source environment. The Emulex NPIV driver is fully open source and has been posted on the sourceforge.net community site for partner evaluation and joint projects. Several such joint projects are in various stages of development.
- Other joint initiatives leveraging NPIV in both virtual and traditional operating systems, and on multiple processor architectures, are under way.

A number of features already available in existing Emulex HBAs and drivers take on added prominence and value in virtual server environments:

- Mixed workload performance: multiple virtual machines typically generate a mixed workload flowing through the same physical adapter(s). By implementing flow control at the frame level, and not just at the sequence level as other vendors do, the Emulex firmware avoids situations where a large block request (e.g., file transfer or database checkpoint) from one of the guest operating systems, monopolizes resources and severely affects response to the smaller requests (e.g., on-line transactions). Frame-level exchanges avoid these situations and ensure the overall throughput is maximized and each application or guest OS is served promptly. The concept of "fairness" implemented through frame-level exchanges maximizes utilization on SAN connectivity and enhances the benefits of server consolidation.
- Predictable Quality of Service: frame-level exchanges introduce fairness in mixed workload by mitigating the impact of background large-block operations on mission-critical transactions, for more predictable OLTP response times and QoS. Frame-level exchange technology, which is built into all Emulex HBA products,

becomes of paramount value in virtual servers as they are designed to support and encourage mixed I/O workloads, especially as they become of age and able to re-host major ported applications and corporate database servers.

- Quality of Service in Virtual Server environments: NPIV provides the technology infrastructure to apply end-to-end application-specific QoS capabilities within virtual server environments.

Summary

Emulex is currently the only vendor offering NPIV support in its HBAs, both traditional and blade-sized, and drivers. NPIV has been implemented cooperatively by Emulex and IBM and is now successfully deployed on the IBM zSeries mainframe product line, an environment where virtual machines are extensively used in demanding production environments. NPIV-capable drivers for Windows, Linux and VMware have been demonstrated and are now in the final stages of productization.

Emulex is capitalizing on its technology advantage, in working with the leading virtualization vendors to architect, deploy and bring to market NPIV-based virtual environments. Important demonstrations, announcements and product introductions will be staged in the coming months; visit www.emulex.com for more information.

Corporate HQ: 3333 Susan Street, Costa Mesa, CA 92626 (714) 662-5600
Wokingham, U.K. (44) 118-977-2929 | **Paris, France** (33) 41-91 19-90
Beijing, China (86 10) 6849-9547

07-344 12/06